Quality Control Under Markovian Deterioration

by

Sheldon Ross

TECHNICAL REPORT NO. 110

June 14, 1968

Supported by the Army, Navy, Air Force and NASA under

Contract Nonr-225(53) (NR-042-002)

with the Office of Naval Research

Gerald J. Lieberman, Project Director

Department of Operations Research

and

Department of Statistics

Stanford University

Stanford, California

Quality Control Under Markovian Deterioration

by

Sheldon Ross

1. Introduction

This paper considers the following model: A production process produces items at the beginning of distinct time periods $t = 0,1,2,\ldots$ It is supposed that at any time t the production process may be in any one of a countable number of states $0,1,2,\ldots$ and that the quality of the item produced is a function of this underlying state. It is also supposed that the state of the process of time t is not known and can only be determined by sampling the item produced. If the process is in state i then a cost I_i is involved in sampling the item. The purpose of sampling is not to replace poor items by good ones but rather to check the manufacturing process.

Thus at the beginning of a period one must decide whether to inspect the item produced or not. Also one may decide to revise the process. This might be done, for instance, if an item had been sampled the previous period and had shown that the production process was in a poor state. The cost associated with revising a process in state i is R_i.

It is supposed that if the process is revised at the beginning of period t then it will be in state 0 at the end of period t. Also it is assumed that no item is produced during that period. If the process is in state i at the beginning of period t and is not revised then it

will remain in state i during the remainder of that period. If a process is in state i at the end of a period then with probability P_{ij} it will be in state j at the beginning of the next period.

If the process in in state i and an item is produced without inspection then there is a cost C_i incurred. (The inspection cost I_i may be thought of as already including a production cost). However we suppose that this cost does not directly become known and thus cannot be used to indicate the state of the process. It is assumed that all costs and transition probabilities are known, and that all costs are bounded.

In this paper a framework is provided for handling problems of this nature. In Section 2 a method of indicating the (observed) "state" of the system at any time t is given and some theorems relating to the convexity of the optimal inspection and revision regions are proven. In section 3 a two-state production process is considered and the structure of the optimal policy is determined. An interesting sidelight of this is that the optimal policy doesn't necessarily have the simple form which intuition might lead one to predict. In section 4 we treat the case where one of the parameters of the model is not fully known.

The general model considered here is similar to one considered in [2]. However, both the methods employed and the results obtained are different. It should also be mentioned that the above model need not be interpreted solely in a quality control context but may also be interpreted as a model for machine deterioration when inspection is costly.

2. General Model

We shall say that the system is in state $P = (P_0, P_1, \ldots)$ if with probability P_i the underlying process is in state i, $i = 0,1,\ldots$ We let $S = \{P = (P_0, P_1, \ldots\ldots): P_i \geq 0, \sum_i P_i = 1\}$ denote the state space of the system. Thus we are allowing for the possibility that at time $t = 0$ we only know the underlying state of the production process up to some arbitrary probability distribution.

Let X_t denote the state of the system at the beginning of period t; and let Δ_t denote the action chosen at t- either produce without inspection (P), produce with inspection (I), or revise the process (R).

$$\text{Let}\quad C(X_t, \Delta_t) = \begin{cases} \sum_i P_i C_i & \text{if } X_t = (P_0, \ldots\ldots) \text{ and } \Delta_t = \underline{P} \\ \sum_i P_i I_i & \text{if } X_t = (P_0, \ldots\ldots) \text{ and } \Delta_t = \underline{I} \\ \sum_i P_i R_i & \text{if } X_t = (P_0, \ldots\ldots) \text{ and } \Delta_t = \underline{R} \end{cases}$$

A policy is any (measurable) rule for choosing actions. For any policy R and $\beta \varepsilon (0,1)$ let $\psi(P, \beta, R) = \sum_{t=0}^{\infty} \beta^t E_R[C(X_t, \Delta_t) \mid X_0 = P]$, and let $V_\beta(P) = \inf_R \psi(P, \beta, R)$ $P \varepsilon S$. Thus $V_\beta(P)$ is the expected cost incurred when an optimal policy is employed given that the system starts in state P and future costs are discounted by a factor β.

For any $P \varepsilon S$, let $TP = ((TP)_0, (TP)_1, \ldots)$ where $(TP)_i = \sum_j P_j P_{ji}$ $i = 0,1,\ldots$ and let $e^i = (P_{i0}, P_{i1}, \ldots)$ $i = 0,1,\ldots$

Then $P\{X_{t+1} = TP \mid X_t = P, \Delta_t = \underline{P}\} = 1$

$$P\{X_{t+1} = e^i \mid X_t = P, \Delta_t = \underline{I}\} = P_i \qquad i = 0,1,\ldots.$$

$$P\{X_{t+1} = e^0 \mid X_t = P, \Delta_t = \underline{R}\} = 1$$

It is well known (see [1]) that $V_\beta(P)$ is the unique solution to

$$(1) \quad V_\beta(P) = \min\{\Sigma P_i C_i + \beta V_\beta(TP);\ \Sigma P_i I_i + \beta \Sigma P_i V_\beta(e^i);$$

$$\Sigma P_i R_i + \beta V_\beta(e^0)\} \quad P \varepsilon S$$

and any rule R_β which when in state P selects an action which minimizes the right side of (1) is β-optimal - i.e. $\psi(P,\beta,R_\beta) = V_\beta(P)$ for all $P\varepsilon S$

Definition: The β-optimal produce region $\equiv \{P:V_\beta(P) = \Sigma P_i C_i + \beta V_\beta(TP)\}$

The β-optimal inspect region $\equiv \{P:V_\beta(P) = \Sigma P_i I_i + \beta \Sigma P_i V_\beta(e^i)\}$

The β-optimal revise region $\equiv \{P:V_\beta(P) = \Sigma P_i R_i + \beta V_\beta(e^0)\}$

Lemma 2.1: $V_\beta(P)$ is a concave function of P - i.e. if $P = \lambda P^1 + (1-\lambda)P^2$ Then $V_\beta(P) \geq \lambda V_\beta(P^1) + (1-\lambda)V_\beta(P^2)$.

Proof: Let $V_\beta^1(P) = \min\{\Sigma P_i C_i, \Sigma P_i I_i, \Sigma P_i R_i\}$

$$V_\beta^n(P) = \min\{\Sigma P_i C_i + \beta V_\beta^{n-1}(TP);\ \Sigma P_i I_i + \beta \Sigma P_i V_\beta^{n-1}(e^i);$$
$$\Sigma P_i R_i + \beta V_\beta^{n-1}(e^0)\}$$

Then $V_\beta^1(P)$ being the minimum of three concave functions is concave. Assuming that $V_\beta^{n-1}(P)$ is concave we get that $V_\beta^n(P)$ is concave for the same reason since $T(\lambda P' + (1-\lambda)P^2) = \lambda T(P') + (1-\lambda)\, T(P^2)$. Thus, by induction, $V_\beta^n(P)$ is concave for all n. But $V_\beta^n(P)$ is just the minimal expected costs incurred over n stages and thus $V_\beta^n(P) \to V_\beta(P)$.

QED.

Theorem 2.2: Both the β-optimal inspect and revise regions are convex.

Proof: Suppose $V_\beta(P^1) = \Sigma P_i^1 C_i + \beta \Sigma P_i^1 V_\beta(e^i)$

and $V_\beta(P^2) = \Sigma P_i^2 C_i + \beta \Sigma P_i^2 V_\beta(e^i)$

and let $P = \lambda P^1 + (1-\lambda)P^2$. Then

$$V_\beta(P) \geq V_\beta(P^1) + (1-\lambda)V_\beta(P^2) = \Sigma P_i C_i + \beta \Sigma P_i V_\beta(e^i)$$

but by (1) we get the reverse inequality. The same method works for the revise region.

QED.

Often one is interested in an optimality criterion which does not discount future costs. Such a criterion is the long-run expected average cost per unit time. So for any policy R define

$\phi(P,R) = \limsup_{n\to\infty} \sum_{t=0}^{n} E_R[C(X_t,\Delta_t)|X_0 = P]/n$. Fix some $P^0 \varepsilon S$ and let

$f_\beta(P) = V_\beta(P) - V_\beta(P^0)$. The following theorem was proven by Ross in [1].

Theorem: If $\{f_\beta(P): P\varepsilon S, \beta\varepsilon(0,1)\}$ is a uniformly bounded equicontinuous family of functions then

(a) There exists a bounded function f(P) and a constant g such that

(2) $g + f(P) = \min\{\Sigma P_i C_i + f(TP);\ \Sigma P_i I_i + \Sigma P_i f(e^i);\ \Sigma P_i R_i + f(e^0)\}\ P\varepsilon S$

(b) $g = \lim_{\beta\to 1} (1-\beta)V_\beta(P)$ for all $P\varepsilon S$; and for some sequence $\beta_r \to 1$

$f(P) = \lim_{r\to\infty} f_{\beta_r}(P)$

(c) If R* is any rule which when in state P selects an action which minimizes the right side of (2) then

$$g = \phi(P,R^*) = \min_R \phi(P,R) \text{ for all } P\varepsilon S$$

From (b) of the above and Lemma 2.1 we thus have

Lemma 2.1': If $\{f_\beta(P)\}$ is uniformly bounded and equicontinuous then $f(P)$ is a concave function of P.

Theorem 2.2': If $\{f_\beta(P)\}$ is uniformly bounded and equicontinuous then both the average-cost optimal inspect and revise regions are convex.

Proof: Same as Proof of Theorem 2.2. (The average-cost regions are defined by using equation (2) in the same manner as equation (1) was used in the β-discount case).

3. A Two-State Production Process

In this section we shall suppose that there are two underlying states - 0 (the good state) and 1 (the bad state). If the process is in the good state at time t and, if the process isn't revised, then with probability π it will be in the bad state at time t+1 where it will remain until it is revised - i.e. $P_{00} = 1 - \pi$, $P_{11} = 1$.

The cost of producing without inspection will be taken to be zero for the good state ($C_0 = 0$) and C for the bad state ($C_1 = C$). The inspect cost I and the revise cost R will be assumed not to depend on the underlying state - i.e. $I_0 = I_1 = I$, $R_0 = R_1 = R$. It shall be assumed throughout that $C < I < R$ (this conditions is natural since the inspect cost is supposed to include some cost due to production).

We will only consider the discounted-cost case and will be concerned with determining the structure of the optimal policy rather than with computational algorithms.

Since there are only two states we may let $S = \{P : P \varepsilon [0,1]\}$; and we say that $X_t = P$ if P is the probability that at the beginning of period t the underlying process is in the bad state. Also in this specialization of our general model we have that $TP = P + \pi - \pi P$ and

$$(3) \qquad V_\beta(P) = \min\{CP + \beta V_\beta(TP);\ I + \beta P V_\beta(1) + \beta(1-P)V_\beta(\pi);\ R + \beta V_\beta(\pi)\} \qquad P\varepsilon[0,1].$$

Lemma 3.1: $V_\beta(P)$ is monotone non-decreasing in P.

Proof: Let

$$(4) \qquad \begin{aligned} &V_\beta^1(P) = \min\{CP, I, R\} \qquad \text{and recursively} \\ &V_\beta^n(P) = \min\{CP + \beta V_\beta^{n-1}(TP);\ I + \beta P V_\beta^{n-1}(1) + \beta(1-P)V_\beta^{n-1}(\pi);\ R + \beta V_\beta^{n-1}(\pi)\} \end{aligned}$$

Then it is easily seen by induction that $V_\beta^n(P)$ is monotone for all n and thus that $V_\beta(P)$ is monotone.

QED.

Lemma 3.2: Every β-optimal policy produces at all P such that $0 \leq P \leq \pi$.

Proof: Suppose some optimal policy inspects at $P\varepsilon[0,\pi]$. Then

$$V_\beta(P) = I + \beta P V_\beta(1) + (1-P)V(\pi) \geq I + \beta V_\beta(\pi) \geq I + \beta V_\beta(P)$$

by monotonicity. Thus $V_\beta(P) \geq I/1-\beta$.

However by (3) $V_\beta(\cdot) \le C + \beta V_\beta(1)$ and thus $C/1-\beta \ge V_\beta(1) \ge V_\beta(P) \ge I/1-\beta$ which is a contradiction. A similar kind of contradiction is arrived at if an optimal policy revises at $P\varepsilon[0,\pi]$.

QED.

<u>Theorem 3.3</u>: An optimal policy R_β may be determined by three numbers P_1,P_2,P_2 $\pi \le P_1 \le P_2 \le P_3 \le 1$ such that R_β produces for $0 \le P < P_1$, inspects for $P_1 \le P < P_2$, produces for $P_2 \le P < P_3$ and revises for $P \ge P_3$.

<u>Proof</u>: By Lemma 3.1 and (3) it follows that the β-optimal revise region may be taken to be a right-hand interval. The result then follows from Lemma 3.2 and Theorem 2.2.

QED.

Thus the β-optimal policy may be described graphically as follows:

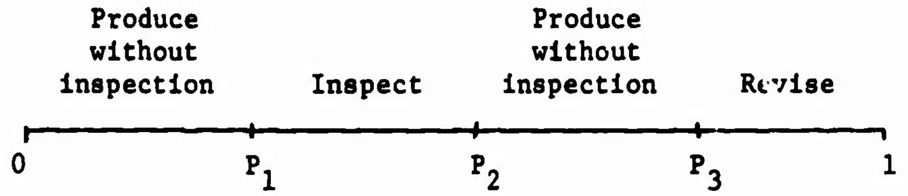

It is however somewhat counter-intuitive to have two disjoint produce regions. Intuitively it would seem likely that the second produce region could always be taken to be vacuous. That this is not so, and thus that sometimes four distinct regions are necessary to characterize the β-optimal policy, is shown by the following example.

__Example__: $C = 4$, $\pi = .1$, $R = 10$, $I = 6$.

Then by letting $\beta = 1$ and using (4) we can show that

$$V_1^7(P) = \begin{cases} 20.87\,P + 7.13 & P \leq .471 \\ 8.13\,P + 13.13 & .471 \leq P < .493 \\ 17.13 & P \geq .493 \end{cases}$$

and thus

$$V_1^8(P) = \begin{cases} \min\{22.78\,P + 9.22;\ 7.91\,P + 15.22;\ 19.22\} & \text{for } P < .412 \\ \min\{11.31\,P + 13.94;\ 7.91\,P + 15.22;\ 19.22\} & \text{for } .412 \leq P < .437 \\ \min\{4\,P + 17.13;\ 7.91\,P + 15.22;\ 19.22\} & \text{for } P \geq .437 \end{cases}$$

thus

$$V_1^8(P) = \begin{cases} 22.78\,P + 9.22 & P < .404 & (\underline{P}) \\ 7.91\,P + 15.22 & .404 \leq P < .489 & (\underline{I}) \\ 4P + 17.13 & .489 \leq P < .521 & (\underline{P}) \\ 19.22 & P \geq .521 & (\underline{R}) \end{cases}$$

Thus for β near 1 the β-optimal eight-stage policy starts off by producing for $P\epsilon[0,.404)$, inspecting for $P\epsilon[.404\ .489)$, producing again for $P\epsilon[.489, .521)$, and revising for $P \geq .521$. Thus we see that four distinct action regions mi t be necessary. The next theorems give sufficient conditions for the optimal policy to have a simpler form than the general one givey by Theorem 3.3. For $n \geq 1$ let $T^nP = T(T^{n-1}P)$ where $T^0P \equiv P$, then

$$\begin{aligned} T^1P &= \pi + (1-\pi)P \\ T^2P &= \pi + (1-\pi)\pi + (1-\pi)^2P \\ T^nP &= \pi + (1-\pi)\pi + \ldots + (1-\pi)^{n-1}\pi + (1-\pi)^nP \\ &= 1 - (1-P)(1-\pi)^n \end{aligned}$$

Let R^0 be the policy that always produces without inspection (always takes action $\underline{P}$).

Then $\psi(\beta,P,R^0) = \sum_{n=0}^{\infty} \beta^n C T^n P = C \sum_{n=0}^{\infty} \beta^n \, (1 - (1-P)(1-\pi)^n)$

$$(5) \qquad = \frac{C}{1-\beta} - \frac{C(1-P)}{1-\beta(1-\pi)}$$

<u>Theorem 3.4</u>: (a) R^0 is β-optimal if and only if $R \geq \frac{C}{1-\beta(1-\pi)}$

(b) If $R < \frac{C}{1-\beta(1-\pi)}$ then every β-optimal policy revises for P near 1.

<u>Proof</u>: If $R \geq \frac{C}{1-\beta(1-\pi)}$ then it can be checked by direct substitution that $\psi(P,\beta,R^0)$ satisfies (3) and thus R^0 is optimal. If R^0 is optimal then by (3) we have that

$$\psi(1,\beta,R^0) \leq R + \beta\psi(\pi,\beta,R^0) \text{ which implies by (5) that}$$

$$\frac{C}{1-\beta} \leq R + \frac{\beta\pi C}{(1-\beta)(1-\beta(1-\pi))}$$

$$\text{or} \qquad R \geq \frac{C}{1-\beta(1-\pi)}$$

To prove (b) we note by (3) that if an optimal policy doesn't revise for P = 1, then $V_\beta(1) = \frac{C}{1-\beta} \leq R + \beta V_\beta(\pi) \leq R + \beta\psi(\pi,\beta,R^0) = R + \frac{\beta\pi C}{(1-\beta)(1-\beta(1-\pi))}$

which implies that $R \geq \frac{C}{1-\beta(1-\pi)}$. The result follows for all P near 1 by the continuity of $V_\beta(P)$. (The continuity of $V_\beta(P)$ is proven in the next Lemma).

QED.

The following Lemma will be needed in the sequel.

<u>Lemma 3.5</u>: $$|V_\beta^n(P_1) - V_\beta^n(P_2)| \leq C|P_1 - P_2| \frac{1 - (\beta(1-\pi))^n}{1 - \beta(1-\pi)}$$

all P_1, P_2, all n.

<u>Proof</u>: The proof is by induction; the result is trivial for $n = 1$. So assume it for $n - 1$. There are now three cases:

(i) $V_\beta^n(P_1) = CP_1 + \beta V_\beta^{n-1}(TP_1)$ which implies by (3) that

$$V_\beta^n(P_2) - V_\beta^n(P_1) \leq C|P_2 - P_1| + \beta[V_\beta^{n-1}(TP_2) - V_\beta^{n-1}(TP_1)]$$

$$\leq C|P_2 - P_1| + \beta C(1-\pi)|P_1 - P_2| \frac{1 - (\beta(1-\pi))^{n-1}}{1 - \beta(1-\pi)}$$

$$= C|P_2 - P_1| \frac{1 - (\beta(1-\pi))^n}{1 - \beta(1-\pi)}$$

(ii) $V_\beta^n(P_1) = I + \beta P_1 V_\beta^{n-1}(1) + \beta(1-P_1) V_\beta^{n-1}(\pi)$ which implies by (3) that

$$V_\beta^n(P_2) - V_\beta^n(P_1) \leq \beta|P_1 - P_2| \quad [V_\beta^{n-1}(1) - V_\beta^{n-1}(\pi)]$$

$$\leq \beta|P_1 - P_2| \quad C(1-\pi) \frac{1 - (\beta(1-\pi))^{n-1}}{1 - \beta(1-\pi)}$$

$$\leq C|P_1 - P_2| \frac{1 - (\beta(1-\pi))^n}{1 - \beta(1-\pi)}$$

(iii) $V_\beta^n(P_1) = R + \beta V_\beta^{n-1}(\pi)$ which implies that

$$V_\beta^n(P_2) - V_\beta^n(P_1) \leq 0.$$

The result then follows by interchanging P_1 and P_2.

QED

The following corollary is immediate

<u>Corollary 3.6</u>: $|V_\beta(P_1) - V_\beta(P_2)| \leq \dfrac{C|P_1-P_2|}{1-\beta(1-\pi)} \leq \dfrac{C|P_1-P_2|}{\pi}$

<u>Theorem 3.7 (Sufficient Conditions)</u>;

(a) $\dfrac{(1-\beta)(R+\beta V_\beta(\pi))}{C} \leq \dfrac{R - I}{\beta(V_\beta(1) - V_\beta(\pi))}$ is a sufficient condition for the existence of a β-optimal policy which produces for $P < P_1$, inspects for $P_1 \leq P < P_2$ and revises for $P \geq P_2$ for some $\pi \leq P_1 \leq P_2 \leq 1$.

(b) $I + \beta(V_\beta(1) - V_\beta(\pi)) \geq \dfrac{C}{1-\beta(1-\pi)}$ <u>or</u> $\dfrac{R - I}{\beta(V_\beta(1) - V_\beta(\pi))} \leq \dfrac{R}{C}(1-\beta(1-\pi))$

is a sufficient condition for the existence of a β-optimal policy which produces for $P < P_1$ and revises for $P \geq P_1$ for some $P_1 \geq \pi$ - i.e. no inspection region.

<u>Proof</u>: (a) Let P_1 and P_2 be such that $CP_1 + \beta V_\beta(TP_1) = R + \beta V_\beta(\pi)$ and $I + \beta P_2 V_\beta(1) + \beta(1-P_2)\, V_\beta(\pi) = R + \beta V_\beta(\pi)$. If such a P_i doesn't exist then let it be infinite, $i = 1,2$. Then using the fact that $CP + \beta V_\beta(TP)$ is monotone and concave it follows that a necessary condition for every β-optimal policy to have four distinct action intervals is for $P_1 > P_2$. (See figure 1).

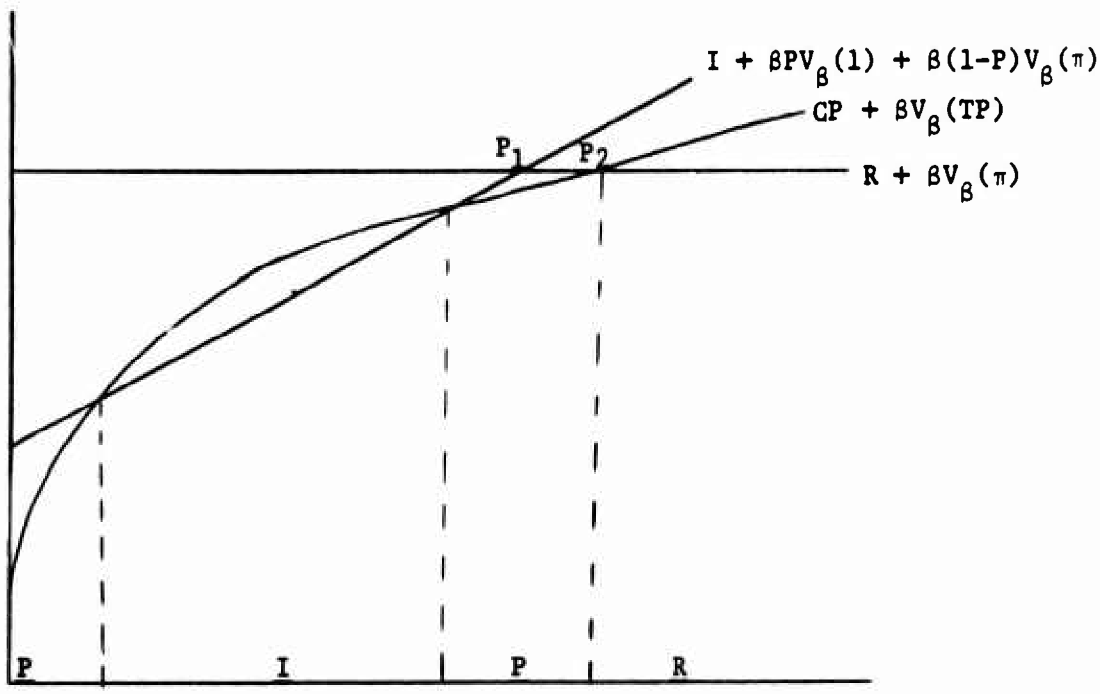

Figure 1

But if $P_1 > P_2$ then $V_\beta(P_1) = R + \beta V_\beta(\pi)$ and thus by monotonicity $V_\beta(TP_1) = R + \beta V_\beta(\pi)$. Thus $CP_1 + \beta(R+\beta V_\beta(\pi)) = R + \beta V_\beta(\pi)$ or

$$P_1 = \frac{(1-\beta)(R+\beta V_\beta(\pi))}{C}.$$ Thus $P_1 > P_2$ implies that

$$(1-\beta)\frac{(R+\beta V_\beta(\pi))}{C} > \frac{R - I}{\beta(V_\beta(1) - V_\beta(\pi))}$$ and thus (a) is proven.

(b) In order for every β-optimal policy to inspect at P we must have

(6) $I + \beta P V_\beta(1) + \beta(1-P)V_\beta(\pi) < R + \beta V_\beta(\pi)$ and

(7) $I + \beta P V_\beta(1) + \beta(1-P)V_\beta(\pi) < CP + \beta V_\beta(TP)$.

Now (6) implies that $P < \dfrac{R - I}{\beta(V_\beta(1) - V_\beta(\pi))}$. From (7) and Corollary 3.6 we get

$$I + \beta P V_\beta(1) + \beta(1-P)V_\beta(\pi) < CP + \beta[V_\beta(\pi) + \frac{C}{1-\beta(1-\pi)}(1-\pi)P]$$

which implies that

$$P > \frac{I}{\frac{C}{1-\beta(1-\pi)} - \beta(V_\beta(1) - V_\beta(\pi))}.$$ Thus we would need both that

$$\frac{R - I}{\beta(V_\beta(1) - V_\beta(\pi))} > \frac{I}{\frac{C}{1-\beta(1-\pi)} - \beta(V_\beta(1) - V_\beta(\pi))}$$ and

$$\frac{I}{\frac{C}{1-\beta(1-\pi)} - \beta(V_\beta(1) - V_\beta(\pi))} < 1.$$ Thus if either of the above inequality doesn't hold then there exists a β-optimal policy which never inspects. It is easy to see that it can be taken to have the desired form.

QED

The conditions given in Theorem 3.7 unfortunately depend on $V_\beta(1)$ and $V_\beta(\pi)$. However, we can prove the following:

<u>Corollary 3.8</u>: If $R < \dfrac{C}{1-\beta(1-\pi)}$ then either

$$I + (R-C) \geq \frac{C}{1-\beta(1-\pi)} \quad \underline{\text{or}} \quad \frac{R - I}{\beta(R-C)} \leq \frac{R}{C}(1-\beta(1-\pi))$$

is a sufficient condition for a β-optimal policy which produces for $P < P_1$ and revises for $P \geq P_1$.

Proof: $R < \frac{C}{1-\beta(1-\pi)}$ implies by Theorem 3.4 that $V_\beta(1) = R + V_\beta(\pi)$.

Thus $V_\beta(1) - V_\beta(\pi) = R + (\beta-1) V_\beta(\pi) \geq R - C$. The result follows from Theorem 3.7.

QED

Note that if $R \geq \frac{C}{1-\beta(1-\pi)}$ then R^0, the policy which always produces (without inspection) is optimal.

4. Unknown π

We have assumed up to this point that all the parameters of the model - C,I,R and π - are known. However, while the cost parameters would probably be known it is quite likely that π will not be known with certainty. We shall now give a method for estimating π from past records of the process; we also show what to do if an apriori distribution for π is known.

Estimation of π

We shall suppose that the past records for the process yield the following sort of data: $(n_1,Z_1),\ldots(n_r,Z_r)$ where n_i denotes the number of periods succeeding the time at which the process was known to be in the good state (either by a revision or by an inspection showing it to be good) until it was next inspected, and Z_i is 1(0) if the inspection showed the process to be good (bad).

Then $P\{Z_i = 1\} = 1 - P\{Z_i = 0\} = (1-\pi)^{n_i}$, and so the probability density of Z_i is given by $P_i(Z_i) = (1-\pi)^{n_i Z_i} (1-(1-\pi)^{n_i})^{1-Z_i}$ $Z_i = 0,1$ and the joint likelihood of all the Z_i's is given by

$$L(Z_1,\ldots Z_r) = (1-\pi)^{\sum_1^r n_i Z_i} \prod_{i=1}^{r} (1-(1-\pi)^{n_i})^{1-Z_i}$$

$$\log L(Z_1,\ldots Z_r) = \sum_1^r n_i Z_i \log (1-\pi) + \sum_{i=1}^{r} (1-Z_i) \log (1-(1-\pi)^{n_i})$$

$$\frac{\partial}{\partial\pi} \log L(Z_1,\ldots,Z_r) = -\sum_1^r \frac{n_i Z_i}{1-\pi} + \sum_{i=1}^{r} (1-Z_i)\, n_i (1-\pi)^{n_i-1}[1-(1-\pi)^{n_i}]^{-1}$$

and so the maximum likelihood $\hat{\pi}$ is given by

$$\hat{\pi} = \begin{cases} 1, & \text{if } Z_i = 0 \text{ for all } i \\ 0, & \text{if } Z_i = 1 \text{ for all } i \\ \text{the solution to } \dfrac{1}{1-\hat{\pi}} \displaystyle\sum_1^r n_i Z_i = \sum_{i=1}^{r} (1-Z_i) n_i (1-\hat{\pi})^{n_i-1}[1-(1-\hat{\pi})^{n_i}]^{-1}, \\ \text{otherwise} \end{cases}$$

Special Case: if $n_i = n$ for all $i = 1,\ldots,r$ then

$$\hat{\pi} = 1 - \sqrt[n]{\sum_1^r Z_i/r}$$

(b) Prior Distribution for π

We suppose that we have an apriori density g_0 - i.e. $P\{\pi \leq x\} = \int_0^x g_0(y)dy \quad 0 \leq x \leq 1$ - and that we are interested in minimizing the expected β-discounted costs.

We shall say that the system is in state $(P(\pi),g)$ at time t - i.e. $X_t = (P(\pi),g)$ - if $P(\pi)$ denotes the probability (possibly as a function of the unknown π) that the process is in the bad state at time t, and if g is the posterior (given everything that has happened up to time t) density of π.

For $t = 0,1,\ldots$ let $X_t = (P^t(\pi), g^t)$. We shall assume that $P^0(\pi)$ is either of the form $P^0(\pi) = P$ or $P^0(\pi) = (1-\pi)P + \pi$ where P is some number in $[0,1]$. Thus $P^0(\pi)$ is monotone non-decreasing in π and from this it follows that $P^t(\pi)$ will be monotone non-decreasing in π. This is so because $P^{t+1}(\pi)$ is either π or 1, or $(1-\pi)P^t(\pi) + \pi$. We can thus let the state space

$$S = \left\{ \begin{array}{l} (P(\pi),g): g \text{ is a probability density on } [0,1], \\ 0 \le P(\pi) \le 1 \text{ for all } \pi \in [0,1],\ P(\pi) \text{ is monotone} \\ \text{non-decreasing in } \pi. \end{array} \right.$$

Letting $V_\beta(P(\pi),g)$ denote the expected β-discounted cost incurred over an infinite time span given that the process starts in state $(P(\pi),g)$ and an optimal policy is employed, we have that

$$(8)\quad V_\beta(P(\pi),g) = \min \left\{ \begin{array}{l} CE_g P(\pi) + \beta V_\beta(TP(\pi),g) \\ I + \beta E_g P(\pi) V_\beta(1, g^1_{P(\pi)}) + \beta(1 - E_g P(\pi)) V_\beta(\pi, g^2_{P(\pi)}) \\ R + \beta V_\beta(\pi, g) \end{array} \right.$$

where

$$TP(\pi) = (1-\pi)\, P(\pi) + \pi$$

$$E_g P(\pi) = \int_0^1 P(\pi)\, g(\pi)\, d\pi$$

$$g^1_{P(\pi)}(x) = \frac{P(x)\, g(x)}{\int_0^1 P(x) g(x) dx} = \frac{P(x)\, g(x)}{E_g P(\pi)}$$

$$g^2_{P(\pi)}(x) = \frac{(1-P(x)) g(x)}{\int_0^1 (1-P(x)) g(x) dx} = \frac{(1-P(x)) g(x)}{1 - E_g P(\pi)},$$

where by $P(x)$ we mean $P(\pi)$ evaluated at $\pi = x$.

As before we may also define

$$V_\beta^1(P(\pi),g) = \min\ \{CE_g P(\pi);\ \ I,\ R\} = CE_g P(\pi)$$

$$\text{(9)}\qquad V_\beta^{n+1}(P(\pi),g) = \min \begin{cases} CE_g P(\pi) + \beta V_\beta^n(TP(\pi),g) \\ I + \beta E_g P(\pi)\ V_\beta^n(1,g^1_{P(\pi)}) + \beta(1-E_g P(\pi))V_\beta^n(\pi,g^2_{P(\pi)}) \\ R + \beta V_\beta^n(\pi,g) \end{cases}$$

Thus the finite stage problem may be solved recursively; and

$$V_\beta^n(P(\pi),g) \to V_\beta(P(\pi),g) \quad \text{as } n \to \infty.$$

In this paper we have only considered the case that the true state of the production process is observable upon inspection of the item produced. However often one would not learn the true state upon inspection but would rather get some additional (not necessarily exhaustive) information about the true state. The first paper dealing with this latter model was that of Girshick and Rubin [3]. They however incorrectly stated that the average cost optimal policy may be characterized by three action regions $\frac{\underline{P}\qquad\quad\underline{I}\qquad\quad\underline{R}\quad}{0\qquad\qquad P\qquad\qquad 1}$. The first counter-example showing the Girshick-Rubin solution to be in error was given by Taylor [6]. Tafeen [5] has recently treated a similar model and has shown that under some restrictions on the information pattern and state space the optimal policy may be characterized by three regions. However his result doesn't hold if the state space is allowed to be the whole interval [0,1]. Future research on the general Girshick-Rubin model (under both an average and discounted cost criterion) is thus needed. It would for example be interesting to know if the optimal policy may be characterized by four regions as in the present paper.

References

[1] Blackwell, David (1965), Discounted Dynamic Programming, Ann. Math. Statist. 36, 226-235.

[2] Elfving, G. (1962), Quality Control for Expensive Items. Technical Report No. 57, Department of Statistics, Stanford University

[3] Girshick, M.A., and Rubin, Herman (1952), A Baye's Approach to a Quality Control Model, Ann. of Math. Statist. 23, p. 114.

[4] Ross, Sheldon M. (1968), Arbitrary State Markovian Decision Processes, Technical Report No. 105, Department of Statistics and Department of Operations Research, Stanford University, (to appear in Ann. of Math. Statist.)

[5] Tafeen, Stephen, (1968), A Markov Decision Process with Costly Inspection, Technical Report No. 107, Department of Operations Research and Department of Statistics, Stanford University.

[6] Taylor, Howard (1965), Markovian Sequential Replacement Processes. Ann. Math. Statist. 37, 871-890.

DOCUMENT CONTROL DATA - R&D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Stanford University Dept. of Operations Research and Dept. of Statistics Stanford, California	2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED 2b. GROUP
3. REPORT TITLE Quality Control Under Markovian Deterioration	
4. DESCRIPTIVE NOTES (Type of report and inclusive dates) Technical Report	
5. AUTHOR(S) (Last name, first name, initial) Ross, Sheldon M.	

6. REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
June 14, 1968	19	6

8a. CONTRACT OR GRANT NO. Nonr-225(53) b. PROJECT NO. NR-042-002 c. d.	9a. ORIGINATOR'S REPORT NUMBER(S) Technical Report No. 110 9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)
10. AVAILABILITY/LIMITATION NOTICES Distribution of this document is unlimited	
11. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY Logistics and Mathematical Statistics Branch Office of Naval Research Washington, D.C. 20360

13. ABSTRACT

We suppose that a production process may be in any one of a countable number of states and that the quality of the item produced is a function of this underlying state. It is also supposed that the state of the process at time t is not known and only becomes known when an item is sampled. We suppose that there are inspection, production, and revision costs and provide a framework for handling problems of this nature. We also treat in detail a 2-state version of this problem.

DD FORM 1 JAN 64 1473

UNCLASSIFIED

Security Classification

14. KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Quality Control Markovian Production Process						

INSTRUCTIONS

1. ORIGINATING ACTIVITY: Enter the name and address of the contractor, subcontractor, grantee, Department of Defense activity or other organization *(corporate author)* issuing the report.

2a. REPORT SECURITY CLASSIFICATION: Enter the overall security classification of the report. Indicate whether "Restricted Data" is included. Marking is to be in accordance with appropriate security regulations.

2b. GROUP: Automatic downgrading is specified in DoD Directive 5200.10 and Armed Forces Industrial Manual. Enter the group number. Also, when applicable, show that optional markings have been used for Group 3 and Group 4 as authorized.

3. REPORT TITLE: Enter the complete report title in all capital letters. Titles in all cases should be unclassified. If a meaningful title cannot be selected without classification, show title classification in all capitals in parenthesis immediately following the title.

4. DESCRIPTIVE NOTES: If appropriate, enter the type of report, e.g., interim, progress, summary, annual, or final. Give the inclusive dates when a specific reporting period is covered.

5. AUTHOR(S): Enter the name(s) of author(s) as shown on or in the report. Enter last name, first name, middle initial. If military, show rank and branch of service. The name of the principal author is an absolute minimum requirement.

6. REPORT DATE: Enter the date of the report as day, month, year; or month, year. If more than one date appears on the report, use date of publication.

7a. TOTAL NUMBER OF PAGES: The total page count should follow normal pagination procedures, i.e., enter the number of pages containing information.

7b. NUMBER OF REFERENCES: Enter the total number of references cited in the report.

8a. CONTRACT OR GRANT NUMBER: If appropriate, enter the applicable number of the contract or grant under which the report was written.

8b, 8c, & 8d. PROJECT NUMBER: Enter the appropriate military department identification, such as project number, subproject number, system numbers, task number, etc.

9a. ORIGINATOR'S REPORT NUMBER(S): Enter the official report number by which the document will be identified and controlled by the originating activity. This number must be unique to this report.

9b. OTHER REPORT NUMBER(S): If the report has been assigned any other report numbers *(either by the originator or by the sponsor)*, also enter this number(s).

10. AVAILABILITY/LIMITATION NOTICES: Enter any limitations on further dissemination of the report, other than those imposed by security classification, using standard statements such as:

(1) "Qualified requesters may obtain copies of this report from DDC."

(2) "Foreign announcement and dissemination of this report by DDC is not authorized."

(3) "U. S. Government agencies may obtain copies of this report directly from DDC. Other qualified DDC users shall request through ________________."

(4) "U. S. military agencies may obtain copies of this report directly from DDC. Other qualified users shall request through ________________."

(5) "All distribution of this report is controlled. Qualified DDC users shall request through ________________."

If the report has been furnished to the Office of Technical Services, Department of Commerce, for sale to the public, indicate this fact and enter the price, if known.

11. SUPPLEMENTARY NOTES: Use for additional explanatory notes.

12. SPONSORING MILITARY ACTIVITY: Enter the name of the departmental project office or laboratory sponsoring *(paying for)* the research and development. Include address.

13. ABSTRACT: Enter an abstract giving a brief and factual summary of the document indicative of the report, even though it may also appear elsewhere in the body of the technical report. If additional space is required, a continuation sheet shall be attached.

It is highly desirable that the abstract of classified reports be unclassified. Each paragraph of the abstract shall end with an indication of the military security classification of the information in the paragraph, represented as *(TS)*, *(S)*, *(C)*, or *(U)*.

There is no limitation on the length of the abstract. However, the suggested length is from 150 to 225 words.

14. KEY WORDS: Key words are technically meaningful terms or short phrases that characterize a report and may be used as index entries for cataloging the report. Key words must be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location, may be used as key words but will be followed by an indication of technical context. The assignment of links, roles, and weights is optional.

UNCLASSIFIED

Security Classification